disk drive, the throughput of the storage system will be that of the single disk drive and the reliability of the subsystem will be that of a particular disk drive. --

Please amend the specification by replacement paragraph of paragraph 3 under the section entitled "Background of the Invention" on page 1, lines 19-28, pursuant to 37 CFR 1.121(b)(1) as follows:

-- Redundant arrays of inexpensive disks ("RAID") storage systems have addressed these needs by providing redundancy for reliability and management techniques to achieve higher performance. Specifically, RAID subsystems apply various management techniques (often referred to as RAID "levels") to provide redundancy in the storage of data on the disk drives such that failure of a single disk drive does not render the entire subsystem unusable. Other RAID techniques ("striping") distribute the data over multiple disk drives to achieve the benefit of multiple disk drives processing a single larger I/O request to read or write data. Where *N* disk drives are used to process a single I/O request, the time to complete the request as compared to a single drive is on the order of 1/*N*. --

Please amend the specification by replacement paragraph of paragraph 4 under the section entitled "Background of the Invention" on page 1, line 29, through page 2, line 6, pursuant to 37 CFR 1.121(b)(1) as follows:

-- The "array" of multiple disk drives in a RAID storage subsystem is managed by a RAID storage controller device. The storage controller typically includes a general purpose microprocessor with associated program memory, cache memory for caching data sent to and from the disk drive array, "back-end" interfaces to adapt the controller to the disk drive array (i.e., SCSI and/or Fibre Channel interface controllers), a "front-end" interface to couple the controller to one or more host systems, etc. The storage controller manages the disk array to make the array appear to a host computer as a large single disk drive that offers improved performance and reliability as compared that of a single disk drive. --

Please amend the specification by replacement paragraph of paragraph 5 under the section entitled "Background of the Invention" on page 2, lines 7-15, pursuant to 37 CFR 1.121(b)(1) as follows:

-- To further enhance reliability and performance, RAID subsystems also are known to utilize multiple such storage controllers. The multiple storage controllers are often configured and managed to provide redundancy such that failure of a single storage controller does not render the subsystem inaccessible. The multiple controllers may also be configured to enhance performance of the storage subsystem by providing parallel processing by multiple controllers of multiple host system I/O requests. The load of I/O requests may therefore be distributed over the plurality of storage controllers to reduce the total processing time required for a series of I/O requests that may be processed in parallel. --

Please amend the specification by replacement paragraph of paragraph 1 under the section entitled "Summary of the Invention" on page 3, lines 18-28, pursuant to 37 CFR 1.121(b)(1) as follows:

-- The present invention solves the above and other problems, thereby advancing the state of the useful arts, by providing a storage subsystem architecture that divides the controller function between front-end controllers and back-end controllers and that applies storage area network ("SAN") techniques and devices within the storage subsystem to interconnect the front-end controllers and back-end controllers. SAN components are known and applied outside the storage subsystem for interconnection of such storage subsystems to host computers and other computing subsystems. In the context of this invention, SAN switches are applied ***within*** the storage subsystem to permit more flexible configuration of front-end and back-end control devices within the storage subsystem. --

Please amend the specification by replacement paragraph of paragraph 3 under the section entitled "Summary of the Invention" on page 4, lines 8-14, pursuant to 37 CFR 1.121(b)(1) as follows:

-- Each FEC and BEC includes a SAN interface to connect to the SAN switches. The SAN switches therefore provide flexible interconnection between virtually any number of front-end controllers and any number of back-end controllers. Such a storage subsystem may thereby be flexibly configured to add additional back-end controllers where required for back-end performance or reliability enhancement and may be configured to add additional front-end controllers when required for front-end performance and reliability. --

Please amend the specification by replacement paragraph of paragraph 3 under the section entitled "Detailed Description of the Preferred Embodiments" on page 5, lines 21-28, pursuant to 37 CFR 1.121(b)(1) as follows:

-- As is known in the art, the host communication media 160 may be any of several well-known media including: parallel SCSI, Fibre Channel, Ethernet (or other local area network media), etc. Similarly, it is known in the art that the back-end communication media 150 may be any of several well-known media including parallel SCSI, Fibre Channel, ATA, EIDE, etc. Those skilled in the art will recognize that, depending upon the choice of media, elements 150 and 160 may include appropriate switches, hubs and other connectivity devices as required for the particular communication medium. --

Please amend the specification by replacement paragraph of paragraph 4 under the section entitled "Detailed Description of the Preferred Embodiments" on page 5, line 29, through page 6, line 9, pursuant to 37 CFR 1.121(b)(1) as follows:

-- This exemplary known architecture provides redundant connectivity within the storage subsystem between the storage controllers and the storage modules. As noted above, this known architecture is inflexible in terms of scalability in that the front-end

control functions (i.e., performed within 102) are integrated on a single controller along with the back-end control functions (i.e., performed within 104 and 106). If the subsystem has a need for enhancing back-end performance, additional back-end performance in the form of back-end interface elements and functions are coupled with front-end control circuits and functions. Likewise, if additional front-end processing power for host generated I/O requests is required, the additional controllers are integrated with potentially extraneous back-end control circuits and functions. Furthermore, the interconnection of additional controllers with existing storage modules may be cumbersome depending on the type of connections used. --

Please amend the specification by replacement paragraph of paragraph 5 under the section entitled "Detailed Description of the Preferred Embodiments" on page 6, line 10, through page 7, line 3, pursuant to 37 CFR 1.121(b)(1) as follows:

-- More specifically, the front-end controllers perform processing related to transactions with attached host computer systems and higher level storage management functions while back-end controllers perform processing related to RAID management of the storage devices and lower level controls within the storage subsystem. Each controller therefore addresses different aspects of the overall performance of the storage subsystem. Both front-end and back-end controllers confront problems with available capacity to handle host I/O transactions. The size and frequency of host I/O requests impact the performance requirements of both the front-end controllers and the back-end controllers. Back-end controllers confront problems relating to interfacing with disk drives and the associated communication channels used therefore. In particular, the back-end controller is matched to a communication channel bandwidth associated with a number of disk drives. The configuration of back-end controllers is therefore preferably matched to the performance characteristics of the disk drives attached to it and the associated communication channel bandwidth. A few high performance disk drives can saturate the communication channels used to communicate with back-end controllers. Additional communication bandwidth for disk drives may therefore require additional back-end controllers to accommodate the potential saturation of the disk interface channel. The

needs to scale the front-end transaction processing performance and high end storage management is largely distinct from the needs to scale the back-end performance for RAID management and lower level storage management functions. Though not enabled by prior techniques, it is useful to isolate these functions to permit independent scaling of the performance of front-end control functions and independent scaling of the back-end control functions. --

Please amend the specification by replacement paragraph of paragraph 3 under the section entitled "Detailed Description of the Preferred Embodiments" on page 13, lines 11-27, pursuant to 37 CFR 1.121(b)(1) as follows:

-- In particular, BEC 260 includes one or more SAN interfaces 262 to connect to the SAN communication media 256. The SAN interfaces 262 are coupled via bus 450 to disk interfaces 400 and 402 which, in turn, are coupled via bus 150 to storage modules and/or individual disk drives. As shown in figure 4, disk interfaces 400 and 402 include all intelligence required to interface with a front-end control element via bus 450 and SAN interface 262. Those skilled in the art will recognize that in particular applications it may be beneficial to implement the FEC and BEC as identical hardware components each implementing its particular designated function. Such identity of the hardware components permits more flexible replacement of spare parts in the subsystem. Further, those skilled in the art will recognize that many of the components in an FEC or BEC may be integrated into higher level integrated circuits incorporating many discrete functions into a VLSI custom circuit. Such design choices are well-known to those skilled in the art. Key to the BEC of the present invention is that it is devoid of front-end functions and associated circuits. Rather, it performs only the back-end functions of low level disk drive command processing. Interfacing with higher level front-end control elements is provided via the SAN interfaces of the BEC. --